

SYSTEM AND METHOD FOR REGISTER RENAMING

Inventors: Trevor Anthony Deosoran
Sanjiv Garg
Kevin Ray Iadonato

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation of Application No. 09/826,211, filed April 5, 2001, now allowed, which is a continuation of Application No. 09/399,000, filed September 17, 1999, now U.S. Patent No. 6,272,617, which is a continuation of Application No. 09/063,059, filed April 21, 1998, now U.S. Patent 6,138,231, which is a continuation of Application No. 08/698,211, filed August 15, 1996, now U.S. Patent 5,809,276, which is a continuation of Application No. 08/478,531, filed June 7, 1995, now U.S. Patent No. 5,590,295, which is a continuation of Application No. 07/997,948, filed December 31, 1992, now abandoned.

[0002] The following are related applications:

- *Superscalar RISC Instruction Scheduling*, Serial No. 08/219,425, filed March 29, 1994, now U.S. Patent No. 5,497,499, which is a file wrapper continuation of Serial No. 07/860,719, filed March 31, 1992, now abandoned;
- *Semiconductor Floor Plan for a Register Renaming Circuit*, Serial No. 07/860,718, filed March 31, 1992, now U.S. Patent No. 5,371,684;
- *System and Method for Retiring Instructions in a Superscalar Microprocessor*, Serial No. 08/481,146, filed June 7, 1995, now U.S. Patent No. 5,826,055, which is a file wrapper continuation of Serial No. 07/877,451, filed May 1, 1992, now abandoned;
- *High-Performance, Superscalar-Based Computer System with Out-of-Order Instruction Execution*, Serial No. 07/817,810, filed January 8, 1992,

now U.S. Patent No. 5,539,911, which is a continuation of Serial No. 07/727,006, filed July 8, 1991, now abandoned;

- *High-Performance, Superscalar-Based Computer System with Out-of-Order Execution and Concurrent Results Distribution*, Serial No. 08/397,016, filed March 1, 1995, now U.S. Patent 5,560,032, which is a file wrapper continuation of Serial No. 07/817,809, filed January 8, 1992, now abandoned, which is a continuation of Application No. 07/727,058, filed July 8, 1991, now abandoned.

The above cited patent documents are incorporated herein by reference.

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0003] The present invention relates to superscalar reduced instruction set computers (RISC), more particularly, the present invention relates to a register renaming circuit for superscalar RISC computers.

2. Background Art

[0004] A more detailed description of some of the basic concepts discussed in this application is found in a number of references, including Mike Johnson, *Superscalar Microprocessor Design* (Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1991) (hereafter Johnson); John L. Hennessy et al., *Computer Architecture - A Quantitative Approach* (Morgan Kaufmann Publishers, Inc., San Mateo, California, 1990). Johnson's text, particularly Chapters 2, 6 and 7 provide an excellent discussion of the register renaming issues addressed by the present invention.

[0005] A major consideration in a superscalar RISC processor is to how to execute multiple instructions in parallel and out-of-order, without incurring data

errors due to dependencies inherent in such execution. Data dependency checking, register renaming and instruction scheduling are integral aspects of the solution. A detailed discussion of storage conflicts, register renaming and dependency mechanisms is found in commonly owned, co-pending U.S. Patent No. 5,497,499 (hereafter referred to as the '499 patent).

[0006] The '499 patent discloses a register renaming circuit (RRC) having a fixed instruction window. In the fixed instruction window, dependency checking and resource assignment is done on the same set of instructions until all the instructions in the set are completed. This means that there is a fixed relationship between each instruction's position in the instruction stream and the instruction number used for resource assignment.

[0007] For example, in an instruction stream with an instruction window of 4 instructions every fourth instruction maps to instruction A (i.e., the first instruction in the window). This technique makes re-mapping of instruction very simple. In this case a 4-to-1 multiplexer is all that is necessary for each resource to forward a single instruction to that resource. However, the fixed format requires that the instruction window be advanced by a fixed amount, which results in somewhat inefficient processing.

[0008] When an instruction retires (an instruction can retire after it has been executed without exception and when all previous instructions have been executed and their results are stored in the register file), its result is moved into a main register file (i.e., the programmable visible register file) and if any instructions were dependent on that instruction, their renamed sources are not needed anymore. In the architecture disclosed in the '499 patent, all instructions' sources are renamed every cycle. This renaming technique requires many comparators for performing the dependency checks. More specifically, the source register addresses of each instruction must be compared to the destination register addresses of all preceding instructions in the instruction window every cycle.

[0009] What is desired is a more efficient register renaming technique requiring less comparators and permitting the processor to execute instructions in parallel and out of order.

SUMMARY OF THE INVENTION

[0010] The present invention is directed to a system and method for performing register renaming of source registers on a per-cycle basis only for new instructions added to the instruction window in that cycle. The present invention thus reduces the total number of dependency check comparators necessary for performing register renaming.

[0011] A preferred embodiment of the present invention comprises storing the instructions in a variable advance instruction window, and assigning a tag to each instruction in the instruction window. The tag of each retired instruction is assigned to the next new instruction to be added to the instruction window. The results of instructions executed by the processor are stored in a temp buffer according to their corresponding tags to avoid output dependencies and anti-dependencies. The temp buffer therefore permits the processor to execute instructions out of order and in parallel.

[0012] Data dependency checks are performed only for each new instruction added to the instruction window. Operands of the instructions having input dependencies are often located in the temporary buffer, and the source register addresses of those instructions having dependencies are renamed according to the tags of the operands located in the temp buffer. The renamed source register addresses are then stored in a rename result register file.

[0013] The foregoing and other features and advantages of the present invention will be apparent from the following more particular description of the preferred embodiments of the invention, as illustrated in the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] The invention will be better understood if reference is made to the accompanying drawings. A brief description of the drawings is as follows:

[0015] FIG. 1 shows a representative block diagram of a DDC equal compare circuit of the present invention.

[0016] FIG. 2 shows a representative block diagram of an N-1 input priority encoder of the present invention.

[0017] FIG. 3 shows a representative block diagram of the tag assignment logic (TAL) of the present invention.

[0018] FIG. 4 shows a representative block diagram of the TAL and priority encoder circuit of the present invention.

[0019] FIG's. 5A and 5B show representative block diagrams of the register rename block of the present invention.

[0020] FIG. 6 shows a representative block diagram of the register rename register file of the present invention.

[0021] FIG. 7 shows a representative block diagram of the interconnection of the blocks of FIG's. 5B and 6.

[0022] FIG. 8 shows a representative high level block diagram including the RRC of the present invention.

[0023] FIG. 9 shows a representative block diagram of a circuit to generate the address for one register file port.

DETAILED DESCRIPTION OF THE INVENTION

[0024] The terms processor, CPU, and digital processor are often used interchangeably in this field. The term "processor" is used hereafter with the understanding that other similar terms could be substituted therefore without changing the underlying meaning of this disclosure.

[0025] The present invention is directed to a Register Renaming Circuit (RRC) which is part of a processor. The RRC permits the processor to execute instructions in parallel and out of order. In a preferred embodiment of the present invention, the processor has a variable advance instruction window (VAIW) for holding instructions from an instruction stream prior to execution. The RRC can be used with a fixed advance instruction window as well.

[0026] The VAIW in a preferred embodiment holds eight instructions, and up to four new instructions can be added to the top four locations of the VAIW in any one cycle.

[0027] In a VAIW, any one of instructions 10, 11, 12 and 13 can be mapped into the first location in the window (location A, for example). Tags are assigned to the instructions as the instructions enter the VAIW. The tags are stored in a first-in-first-out buffer (hereafter called a FIFO; not shown).

[0028] As an instruction advances in the VAIW by a variable amount, the tag associated with that instruction also advances in the FIFO by the same amount. When a new instruction enters the VAIW it is assigned the tag of the most recent instruction to leave the VAIW, thus tags are reused. Instructions can leave the VAIW by either retiring, or they can be flushed out if a branch is taken.

[0029] The tag of each instruction leaving the instruction window is returned to the head of the FIFO and re-used by the new instruction added to the window. However, the first instruction and tag in the FIFO always progress in order, because instructions always retire in order.

[0030] According to the present invention only new instructions in the VAIW need be checked for dependencies. This eliminates the need for excess

comparators. New instructions in the window are therefore passed through the RRC. In order to reduce complexity by renaming the source registers for only those instructions that are new in the instruction window on a per-cycle basis, two assumptions are made:

1. Each instruction's tag remains constant as long as the instruction remains in the window. This tag is also associated with the location in a temp buffer (discussed below) that the corresponding instruction's output will be stored.
2. At most, only a subset (I_n to I_{n-1}) of the instructions in the window (I_n to I_0) can be new in any given cycle.

[0031] In a preferred embodiment of the present invention, the temp buffer (or temporary buffer) is part of the main register file. The register file contains 40 registers; registers 0-31 are the main registers (commonly called the programmer visible registers), and registers 32-39 comprise the temp buffer. The temp buffer and main register file share the same read ports. Thus, to read the data from temp buffer address 4, {100100} would be the address on the read address port, for example.

[0032] In order to perform dependency checking in the present invention, an RRC Inputs Register File (RIRF) is used to hold the source and destination register addresses of all instructions in the instruction window. As new instructions enter the window an instruction fetch unit (IFU: not shown) sends the instructions' source and destination register addresses to the RIRF. The source and destination register addresses are stored in the RIRF by tag number. The RIRF has one output for each instruction in the window and the source and destination register addresses are read from the RIRF and sent to the RRC.

[0033] The RRC performs the data dependency checking and the register renaming. The data dependency check is done by the Data Dependency Checker (DDC) and the rename is performed by the Tag Assignment Logic (TAL).

[0034] There are three kinds of data dependencies: *input dependencies*, *anti-dependencies* and *output dependencies*. An instruction is input dependent on a

previous instruction if one of its inputs is the previous instruction's output; an instruction is anti-dependent if the address of one of its source registers (RS) is the same as the address of a later instruction's destination register (RD); and an instruction is output dependent if the address of its destination register is the same as the address of another instruction's destination register.

[0035] Dependencies limit the amount of parallelism that a computer can exploit. For example, if instruction A is input dependent on instruction B, then instruction A must not be executed until after instruction B is executed. Also, if instruction A and B are output dependent and instruction A comes first in the program, then instruction A's result must be written to the main register file before instruction B's. Finally, if instruction A is anti-dependent on instruction B, then instruction B's result must not be written to the main register file until after instruction A has begun executing. Output and anti-dependencies are usually avoided by use of the temp buffer. Input dependencies cannot be avoided and are located by the DDC.

[0036] The DDC locates input dependencies by comparing the register file addresses of each instruction's sources with the register file addresses of each previous instruction's destination. If an instruction's input data comes from the same register file address as a previous instruction's output data, then they are dependent (the term "dependent" will be used to mean "input dependent" for the remainder of this description).

[0037] It is possible that an instruction can be dependent on several previous instructions. When this happens, the RRC assumes that the programmer intended that the instruction be dependent on the most previous instruction. For example, if instruction 5 depends on instructions 3 and 1, then the RRC would assume that the programmer intended instruction 5 to use instruction 3's results and not instruction 1's.

[0038] A DDC equal compare circuit 100 for checking dependencies between instructions A and B in accordance with the present invention is shown in FIG. 1. The output ($A=B?$) of the DDC equal compare circuit 100 is sent to a priority encoder. A $n-1$ input priority encoder 200 is shown in FIG. 2. Priority encoder

200 checks the highest priority dependency for instruction n's source (I_nRS). The inputs at the top of priority encoder 200 are the data dependency comparisons of instruction n's source (I_nRS) with the destinations of all previous instructions ($I_{n-1}RD-I_0RD$), as determined by equal compare circuits 202. For example, if the x^{th} bit of the priority encoder output is asserted, then instruction n is input dependent on instruction x.

[0039] The present invention avoids the problems caused by output and anti-dependencies by storing the results of all instructions in the temp buffer and then moving the results into the main register file in program order. For example, if instruction 1 finishes before instruction 0, its result will be written to the register file after instruction 0's result is written to the register file. The use of the temp buffer allows the processor to execute instructions out of order and in parallel. Since the results of the instructions are moved to the main register file in order, output and anti-dependencies do not cause a problem.

[0040] Since the result of an instruction can sit in the temp buffer for a long time before it gets moved to the main register file, instructions that are dependent on that instruction will also have to wait. In order to improve performance, the present invention includes a means with which to use data that is in the temp buffer. That means is the Tag Assignment Logic (TAL).

[0041] The TAL determines the location in the temp buffer of the operands of dependent instructions. As noted above all instructions are assigned a tag that remains constant while the instruction is in the window, and there is one location in the temp buffer for each instruction in the window. Thus, the processor implementing the present invention uses the tag of an instruction as the temp buffer address of that instruction's result.

[0042] Since the TAL knows where every instruction's result is stored, and since it also knows (from tile DDC) where the dependencies are between instructions, tile TAL can determine the location in the temp buffer of each instruction's inputs.

[0043] A representative block diagram of a TAL 300 used to determine the location of instruction n's source (RS) is shown in FIG. 3. The outputs of the

priority encoder are connected as select lines (as shown generally at 302) to select the I_{n-1} through I_0 TAGs, which are input at the top of TAL 300. TAL 300 thus outputs the temp buffer address of instruction n's source.

[0044] A complete rename circuit for instruction n's source register is shown in FIG. 4. The term for the renamed register file address of instruction n's source is *INRS_TAG*.

[0045] A representative block diagram of a rename circuit 500 of the present embodiment is shown in FIG. 5A. The address of a new instruction's source register (I_nRS) is shown input at the top of the rename circuit 500. The destination register addresses of all preceding instructions in the window are input to rename circuit 500, as shown generally at 502. In addition, all the tags of all preceding instructions in the window are input to rename circuit 500, as shown generally at 504. Rename circuit 500 outputs a tag for the new instruction's source register (I_nRS), as shown at 506. The new I_nRS tag is assigned by rename circuit 500 according to any dependencies, as discussed above in connection with the other embodiment(s). If the instruction has no dependencies the address input at the top input is simply passed to the output. Given a VAIW of 8 instructions and assuming that the temp buffers have the 8 highest addresses of the 40 total registers, the most significant bit of the rename circuit 500 output indicates whether the result is in the main register file or the temp buffer.

[0046] Renaming circuits 508 for renaming the source registers of a plurality of instructions are shown in FIG. 5B. The number "i" in FIG. 5B corresponds to the maximum number of instructions in the window that can be new. In a preferred embodiment "i" is equal to 3; thus requiring 4 rename circuits.

[0047] Because renaming only happens when a new instruction enters the window, some technique is needed to store the outputs of the RRC. One technique would be to store the RRC results in a FIFO whenever the instruction window advances. The present embodiment, however, stores the renamed registers in a separate register file (discussed below). Since each instruction's tag stays constant, the renamed source register results from the rename circuits can be stored by tag

number. The register file for storing the rename results therefore has one write port for each new instruction in the window and one read port for each instruction in the window.

[0048] A representative block diagram of a rename result register file (RRRF) 600 is shown in FIG. 6. The renamed results are input as "WRITE DATA", as shown generally at 601. "READ ADDRESS" and "WRITE ADDRESS" tags are input to register file 600, as shown generally at 602 and 604, respectively. Renamed results for all the sources of all instructions in the window are available as "READ DATA", as shown generally at 606. FIG. 7 shows rename circuits 508 connected to rename result register file 600.

[0049] When an instruction retires its result is moved into the main register file. If any instructions were dependent on that instruction, their renamed sources are not needed anymore.

[0050] The area to which new instructions can enter the instruction window (in this embodiment the area is the top four locations), are those locations which are register renamed. Once an instruction leaves that area of the window it is no longer renamed. The RRC of the present invention renames an instruction's source register when it enters the window, so there needs to be a mechanism to detect which instructions' sources have been moved to the register file and to replace the renamed source register address with the original source register address. The first part of that mechanism is called MONDEP (as in "monitor dependencies") and the second part is called RFMXING. In addition, a source register ready generator (RDY_GEN) is used to determine when each instruction's sources are available.

[0051] A representative high level block diagram of the RIRF, RRC, RRRF, MONDEP, RDY_GEN and RFMXING (labeled as 802, 804, 600, 806, 808 and 810, respectively) is shown in FIG. 8. Each block 802, 804, 600, 806, 808 and 810 receives the tags of all instructions in the instruction window from the tag FIFO (not shown). Implementation of the tag FIFO will become obvious to one skilled in the art.

[0052] Source and destination register addresses of new instructions from the IFU (not shown) are sent to RIRF 802 via a bus 812, and are accessed by RRC 804 via a bus 814. The source registers of all instructions are passed to RFMXING 806 via a bus 816. Renamed source registers of all instructions are stored in RRRF 600 via a bus 818. The stored renamed source registers of all instructions are passed to RFMXING 806, MONDEP 808 and RDY_GEN 810 via a bus 820.

[0053] MONDEP 808 determines which dependencies have disappeared by comparing the tags of retiring or recently-retired instructions with the lower three bits of the renamed sources of each instruction. Information regarding retired instructions is sent to MONDEP 808 via a bus 828 from a retirement unit (not shown; the details of a retirement unit that can be used to generate these signals is disclosed in co-pending, commonly owned U.S. Patent No. 5,826,055). If there is a match, then MONDEP 808 knows that the dependency has been removed and the outputs of MONDEP 808 outputs which instructions' inputs have been moved from the temp, buffer to the register file. These output signals are sent to RFMXING 806 and RDY_GEN 810 via buses 822.

[0054] In a preferred embodiment of the present invention, the instruction window holds eight instructions. Each cycle, at most three of those instructions can be retired. In the cycle after an instruction is retired, its tag is moved to the top of the FIFO. Therefore, to check what dependencies have been removed, MONDEP 808 compares each of the renamed sources of each instruction with the tags of the top three instructions in the FIFO. In a further embodiment MONDEP 808 can compare each renamed source with the tags of the instructions at the bottom of the FIFO that are about to be retired.

[0055] MONDEP 808 outputs a bit for each source of each instruction and the bits are sent to RFMXING and RDY_GEN blocks in the RRC. These bits are asserted when an instruction's dependency goes away and remain asserted until the instruction is retired.

[0056] RDY_GEN 810 determines when each instruction's sources are available and outputs this information via a bus 830. The difference between MONDEP 808

and RDY_GEN 810 is MONDEP 808 only monitors when instructions retire. An instruction does not have to wait until another instruction retires to use its result; it only needs to wait until it is done (an instruction is done when its result is entered into the temp buffer). Also, if an instruction has no dependencies, then it can be executed immediately. Information concerning whether an instruction is "done" is input to RDY_GEN 810 via a bus 832. "Done" signals come from done control logic (not shown). In connection with the present invention, the term "done" means the result of the instruction is in a temporary buffer or otherwise available at the output of a functional unit. (An example of done control logic may be found in the '499 patent.)

[0057] RDY_GEN 810 has one output for each source of all instructions in the window. The output for a particular instruction's source is assured if one of three conditions is true:

1. The source was never dependent on any other instruction.
2. The instruction that the source was dependent on is done and its result is in the temp buffer.
3. The instruction that the source was dependent on is retired and its result has been moved from the temp buffer to the register file.

[0058] These outputs 830 of RDY_GEN 810 go to the ISSUER which determines which instruction(s) is to issued based on functional unit availability and lack of dependencies.

[0059] RFMXING 806 is used to generate the read addresses of the register files. It contains a collection of muxes for each read port of each register file. These muxes are selected by the outputs of the ISSUER and MONDEP 808. Read addresses for each port of each register file are output by RFMXING 806, via a bus 824. (A processor may have a separate register file for a floating point subprocessor and an integer subprocessor, for example.)

[0060] The circuit to generate the address for one register file port is shown in FIG. 9. The ISSUER decides which instructions to execute and which register file ports to use for each instruction by sending select signals via a bus 826 to

RFMXING 806. MONDEP 808 decides which instructions sources have been moved from the register file and which are still inside the temp buffer via bus 822. For example, if one assumes that the ISSUER decides to execute instruction I_n and I_n is dependent on I_{n-1} ; then the ISSUER will select (via select signals 826) two top multiplexers (mux) 902 and 904 and choose I_nRS (I_n 's original source register address) and I_nRS_TAG . If it has, MONDEP 808 will select, using a third mux 906, the output of mux 902 on the left and send I_nRS to the register file. If not, it will choose the output of mux 904 on the right and send the I_nRS_TAG to the register file.

[0061] While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example, and not limitation. Thus the breadth and scope of the present invention should not be limited by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents. All cited patent documents and publications in the above description are incorporated herein by reference.